Metrics for matrix-valued measures via test functions

Lipeng Ning and Tryphon T. Georgiou

arXiv:1409.4097v1 [cs.SY] 14 Sep 2014

Abstract

It is perhaps not widely recognized that certain common notions of distance between probability measures have an alternative dual interpretation which compares corresponding functionals against suitable families of test functions. This dual viewpoint extends in a straightforward manner to suggest metrics between matrix-valued measures. Our main interest has been in developing weakly-continuous metrics that are suitable for comparing matrix-valued power spectral density functions. To this end, and following the suggested recipe of utilizing suitable families of test functions, we develop a weakly-continuous metric that is analogous to the Wasserstein metric and applies to matrix-valued densities. We use a numerical example to compare this metric to certain standard alternatives including a different version of a matricial Wasserstein metric developed in [1], [2].

I. Introduction

Consider the set of probability measures

$$\mathcal{P}(\mathbb{I}) := \{\mu : d\mu(\theta) \geq 0 \text{ for } \theta \in \mathbb{I} \text{ and } \mu(\mathbb{I}) = 1\}$$

where, herein, $\mathbb{I}$ is always thought to be an interval. There is a large family of metrics, comparing $\mu_1, \mu_2 \in \mathcal{P}$, that are expressed in the form

$$\sup_{f \in \mathcal{F}} \left| \int f(d\mu_1 - d\mu_2) \right| \tag{1}$$

with $\mathcal{F}$ being a suitable set of functions on $\mathbb{I}$. Probability metrics that can be expressed in this form (1) are often referred to either as metrics with a ζ-structure [3] or as integral probability metrics [4], [5].

The family of metrics that can be expressed as in (1) includes many familiar ones such as the total variance, Kolmogorov's distance and the 1-Wasserstein metric. In point of fact, the total variation which is defined by

$$\|\mu_1 - \mu_2\|_{\mathrm{TV}} := \int_{\mathbb{I}} |d\mu_1(x) - d\mu_2(x)|,$$

L. Ning is with Brigham and Women's Hospital, Harvard Medical School, Boston, MA 02115, lning@bwh.harvard.edu

T.T. Georgiou is with the Department of Electrical & Computer Engineering, University of Minnesota, Minneapolis, MN 55455, tryphon@umn.edu

The research was supported in part by the NSF under Grant 1027696, the AFOSR under Grant FA9550-12-1-0319, and the Vincentine Hermes-Luh Endowment.

can also be expressed as

$$\|\mu_1 - \mu_2\|_{\mathrm{TV}} = \sup_f \left\{ \int f(d\mu_1 - d\mu_2) \mid \|f\|_\infty \leq 1 \right\},$$

Kolmogorov's metric which is defined by

$$\mathrm{d_K}(\mu_1, \mu_2) := \sup_{x \in \mathbb{I}} |F_1(x) - F_2(x)|$$

with F_1 and F_2 being the cumulative distribution function of μ_1 and μ_2, respectively, can be expressed as

$$\mathrm{d_K}(\mu_1, \mu_2) = \sup_f \left\{ \int_{\mathbb{I}} f(d\mu_1 - d\mu_2) \mid \int_{\mathbb{I}} |f'| dx \leq 1 \right\}$$

see [6, page 73], and so does the 1-Wasserstein metric [7].

We focus on the 1-Wasserstein metric which can be defined over more general spaces. It is defined with respect to the metric $d(x, y) = |x - y|$, as follows. If

$$\mathcal{P}_{\mu_1, \mu_2} \subset \mathcal{P}(\mathbb{I} \times \mathbb{I})$$

denotes the subset of probability measures on $\mathbb{I} \times \mathbb{I}$ that have μ_1 and μ_2 as marginals, the 1-Wasserstein distance between μ_1 and μ_2 is

$$\mathrm{d}_{W_1}(\mu_1, \mu_2) := \inf_{m \in \mathcal{P}_{\mu_1, \mu_2}} \int_{\mathbb{I} \times \mathbb{I}} d(x, y) dm(x, y).$$

Naturally, it can be expressed in a dual form as

$$\mathrm{d}_{W_1}(\mu_1, \mu_2) = \sup_f \left\{ \int_{\mathbb{I}} f(d\mu_1 - d\mu_2) \mid \|f\|_{\mathrm{Lip}} \leq 1 \right\} \tag{2}$$

where

$$\|f\|_{\mathrm{Lip}} := \sup_{x, y \in \mathbb{I}} \frac{|f(x) - f(y)|}{|x - y|}$$

denotes the Lipschitz semi-norm.

An important property of the Wasserstein metric is that it metrizes weak* convergence of probability measures [7, page 212], that is, given μ and any sequence of measures $\{\mu_k, k = 1, 2, \ldots\}$,

$$\mathrm{d}_{W_1}(\mu_k, \mu) \underset{k \to \infty}{\longrightarrow} 0$$

if and only if, for all continuous and bounded f,

$$\int_{\mathbb{I}} f d\mu_k \underset{k \to \infty}{\longrightarrow} \int_{\mathbb{I}} f d\mu.$$

Intuitively, small changes in a weak*-sense reflect small changes in any relevant statistics. This is clearly a desirable feature in any experimental engineering quantification of distances between distributions, whether these represent probability, power, spectral power

or other entities. For this precise reason, the Wasserstein metric has turned out to be a useful tool in modeling of slowly varying time-series [8] and for comparing covariance matrices [9], among many other applications [6].

In the sequel we say that a metric is *weakly continuous* if it metrizes weak* convergence. The use of the terminology "weak" instead of "weak*" is common in probability theory and is adapted in this paper.

The main contribution of the present paper is a particular generalization of the Wasserstein metric to the space of matrix-valued measures with possibly non-equal mass. This metric differs from an analogous metric in our recent work [1] which also represents a generalization of Wasserstein distances to matricial measures. The present metric is based on a dual formalism where we compare measures on a suitable set of test functions. Before dealing with the matricial case, in Section II, we first discuss how to modify the Wasserstein metric so as to compare measures with non-equal masses; a subsequent matricial generalization follows along similar lines. In Section III, we present certain related ideas from non-commutative geometry for devising metrics to compare states of non-commutative algebras; such states are "non-commutative" generalizations of probability measures. In Section IV, we develop the sought weakly-continuous metric between matrix-valued measures. Our interest is in spectral analysis of multivariate time-series and, thereby, we appeal to a pertinent numerical example to highlight differences and similarities of the proposed metric to alternatives; this is given in Section V.

The notational convention we follow is to use regular font, as in μ, m, for scalar values, variables, and functions, and to use boldface fonts, as in $\boldsymbol{\mu}, \boldsymbol{m}$, for matrix-valued functions or elements of a general algebra.

II. Wasserstein-like metric for unbalanced measures

We begin by discussing a certain adaptation of the Wasserstein metric to use on unbalanced measures [10], that is, for the case when the measures we deal with may have unequal integrals. Our motivation stems from the need for a weakly continuous distance to be used on power spectral densities of stationary stochastic processes. The dual formulation serves as a template for a subsequent matricial version.

Consider the set of non-negative scalar measures

$$\mathcal{M}(\mathbb{I}) := \{\mu : d\mu(\theta) \geq 0 \text{ for } \theta \in \mathbb{I}\}.$$

For any two $\mu_1, \mu_2 \in \mathcal{M}$, let

$$\mathrm{d}_{\mathrm{W}_1,\kappa}(\mu_1, \mu_2) := \inf_{\hat{\mu}_1, \hat{\mu}_2} \left\{ \mathrm{d}_{W_1}(\hat{\mu}_1, \hat{\mu}_2) + \kappa \sum_{k=1}^{2} \|\mu_k - \hat{\mu}_k\|_{\mathrm{TV}} \right\} \tag{3}$$

where $\kappa > 0$ is used to weigh in the relative importance of the two terms. In [10], the optimizing variables $\hat{\mu}_1, \hat{\mu}_2$ represent "noise-free" measures having equal mass while the

"error" differences $\mu_k - \hat{\mu}_k$ for $k = 1, 2$, are attributed to statistical variability. The dual of (3) is

$$\mathrm{d}_{\mathrm{W}_1,\kappa}(\mu_1, \mu_2) = \sup_f \{ \int_{\mathbb{I}} f(d\mu_1 - d\mu_2) \mid \|f\|_{\mathrm{Lip}} \leq 1, \|f\|_\infty \leq \kappa \} \tag{4}$$

where $\|f\|_\infty = \max_{x \in \mathbb{I}} |f(x)|$. Since test functions in (4) are bounded, $\mathrm{d}_{\mathrm{W}_1,\kappa}(\mu_1, \mu_2)$ is bounded as well [10], thereby, it can be easily shown that $\mathrm{d}_{\mathrm{W}_1,\kappa}$ is a weakly continuous metric (see [10] for details). Applications of this metric to power spectral analysis has been pursued in [11].

We remark that the essence in (1) is to postulate a family of test functions which is rich enough so as to distinguish measures while the functions in the class be equicontinuous and uniformly bounded [11]. For the Wasserstein metric, the family of test function is the class of Lipschitz functions. A very similar rationale has been introduced in non-commutative geometry, where via a suitable generalization of the Lipschitz semi-norm, one obtains a metric between non-commutative states, namely the Connes' spectral distance. This is discussed next.

III. Metrics in non-commutative geometry

We specialize our discussion to the algebra of $n \times n$ complex-valued matrices $M_n(\mathbb{C})$; for the more general setting of non-commutative algebras see [12], [13]. Given $A, B \in M_n(\mathbb{C})$, or any algebra, we denote by $[A, B] = AB - BA$ the commutator of A and B.

In general, a state $\boldsymbol{\rho}$ of a non-commutative algebra $\mathcal{A}$ is a positive linear functional that maps $f \in \mathcal{A}$ to $\mathbb{R}$ or $\mathbb{C}$. The set of states of $\mathcal{A}$ is denoted as $\mathcal{S}(\mathcal{A})$. For example, a state $\boldsymbol{\rho}$ of the commutative algebra $\mathcal{C}(\mathbb{I})$ of continuous functions on $\mathbb{I}$ uniquely corresponds to a probability measure μ such that for any $f \in \mathcal{C}(\mathbb{I})$

$$\boldsymbol{\rho}(f) = \int_{\mathbb{I}} f d\mu.$$

In quantum mechanics, a state $\boldsymbol{\rho}$ of a non-commutative algebra of "observables" is a *density matrix*; this is positive semidefinite with trace equals to one. For $\boldsymbol{f} \in \mathcal{A}$,

$$\boldsymbol{\rho}(\boldsymbol{f}) = \mathrm{tr}(\boldsymbol{\rho}\boldsymbol{f}).$$

Thus, states are thought of as a generalization of probability measures and $\boldsymbol{\rho}(\cdot) = E_{\boldsymbol{\rho}}\{\cdot\}$ is thought of as the expectation operator.

A. Connes' spectral distance

Consider a non-commutative algebra of operators $\mathcal{A}$ on a Hilbert space. Let D be a specific fixed operator often referred to as the Dirac operator. Connes' spectral distance [12] between $\boldsymbol{\rho}_1, \boldsymbol{\rho}_2 \in \mathcal{S}(\mathcal{A})$ is defined as

$$\mathrm{d}_D(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2) := \sup_{\boldsymbol{f} \in \mathcal{A}} \{ |\boldsymbol{\rho}_1(\boldsymbol{f}) - \boldsymbol{\rho}_2(\boldsymbol{f})| \mid \|[D, \boldsymbol{f}]\| \leq 1 \}.$$

The operator norm of the commutator $[D, \boldsymbol{f}]$ takes the role of the Lipschitz semi-norm. To illustrate the insight in viewing $\|[D, \boldsymbol{f}]\|$ as the analogue of a Lipschitz semi-norm, consider f to be a smooth function on $\mathbb{R}$ and choose $D = -i\partial_x$. For any smooth function g,

$$\begin{aligned}[D, f]g &= Dfg - fDg \\ &= -i\partial_x(fg) + if\partial_x g \\ &= -i\partial_x fg.\end{aligned}$$

Thus $[D, f]$ takes $g \mapsto -i\partial_x fg$. The operator norm $\|[D, f]\|$ equals to the operator norm the mapping $g \mapsto -i\partial_x fg$, which is precisely the Lipschitz semi-norm of f.

Consider $\mathcal{A}$ to be the algebra $\mathcal{C}(\mathbb{I})$ of continuous functions on $\mathbb{I}$. Thus, for any two $\boldsymbol{\rho}_1, \boldsymbol{\rho}_2 \in \mathcal{S}(\mathcal{C}(\mathbb{I}))$ which correspond to probability measures μ_1 and μ_2, respectively,

$$\mathrm{d}_D(\boldsymbol{\rho}_0, \boldsymbol{\rho}_1) = \sup_{\boldsymbol{f} \in \mathcal{A}} \left\{ \int_{\mathbb{I}} f(d\mu_1 - d\mu_2) \mid \|[D, f]\| \leq 1 \right\}.$$

The relation of this distance to the 1-Wasserstein is explained in [13], [14]. We elaborate with an example and provide directions for further generalization.

Connes' spectral distance can be unbounded. For instance, let $\mathcal{A} = M_2(\mathbb{R})$ and let

$$D = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

Consider as states the set $\mathcal{S}(\mathcal{A})$ of positive semi-definite matrices with trace equal to one —observables can also be taken in $\mathcal{A}$. Take an observable

$$\boldsymbol{f} = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

and note that

$$[D, \boldsymbol{f}] = \begin{bmatrix} c - b & d - a \\ a - d & b - c \end{bmatrix}.$$

Thus, the distance between states

$$\boldsymbol{\rho}_0 = \begin{bmatrix} p_0 & q_0 \\ q_0 & 1 - p_0 \end{bmatrix} \text{ and } \boldsymbol{\rho}_1 = \begin{bmatrix} p_1 & q_1 \\ q_1 & 1 - p_1 \end{bmatrix},$$

namely,

$$\begin{aligned} & d_D(\boldsymbol{\rho}_0, \boldsymbol{\rho}_1) \\ & = \sup_{\boldsymbol{f} \in \mathcal{A}} \{ |\mathrm{tr}(\boldsymbol{\rho}_0 \boldsymbol{f}) - \mathrm{tr}(\boldsymbol{\rho}_1 \boldsymbol{f})| \ \mid \ \|[D, \boldsymbol{f}]\| \leq 1 \} \\ & = \sup \left\{ |(p_0 - p_1)(a - d) + (q_0 - q_1)(b + c)| \right. \\ & \qquad \left. \mid \ \| \begin{bmatrix} c - b & d - a \\ a - d & b - c \end{bmatrix} \| \leq 1 \right\} \end{aligned}$$

is ∞. To see this note that if $q_0 \neq q_1$, then $d_D(\boldsymbol{\rho}_0, \boldsymbol{\rho}_1)$ can take any positive value for a suitable choice of $b = c$.

Remark 1: In analogy with (4),

$$\mathrm{d}_{D,\kappa}(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2) := \sup_{\boldsymbol{f}\in\mathcal{A}} \left\{ |\boldsymbol{\rho}_1(\boldsymbol{f}) - \boldsymbol{\rho}_2(\boldsymbol{f})| \mid \|[D, \boldsymbol{f}]\| \leq 1, \|\boldsymbol{f}\| \leq \kappa \right\} \tag{5}$$

defines a bounded metric. Likewise, the Connes' spectral distance can be readily generalized to

$$\mathrm{d}_D(\boldsymbol{\rho}_0, \boldsymbol{\rho}_1) = \sup_{\boldsymbol{f}\in\mathcal{A}} \Big\{ \ |\boldsymbol{\rho}_1(\boldsymbol{f}) - \boldsymbol{\rho}_2(\boldsymbol{f})| \mid \|[D_i, \boldsymbol{f}]\| \leq 1 \quad \text{for } i \in \{1, 2 \ldots, n\} \Big\} \tag{6}$$

for a suitable set of n Dirac operators quantifying "slope" in several possible "directions."

In the next section we deal in more detail with the algebra of matrix-valued continuous functions on $\mathbb{I}$, namely, $\mathcal{C}(\mathbb{C}^{n\times n}, \mathbb{I})$ where states correspond to matrix-valued measures. Our interest stems from spectral analysis of multivariate time-series and in the next section, we present a Wasserstein-like metric between matrix-valued measures. The formalism is completely analogous in that the metric is constructed in a dual formalism by quantifying how measures act on suitably constrained matrix-functions.

IV. A Wasserstein-like metric between matrix-valued measures

Spectral analysis of time-series aims at detecting power and correlations between signals at different parts of the frequency spectrum. Power spectral estimates are typically based on moments, i.e., integrals of the power density, or simply measurements. Hence, it is unreasonable to utilize metrics between power spectral densities that are not weakly continuous.

For the case of multivariable time-series, power densities are matrix-valued. Thus, any metric must weigh in both the frequency content of the power as well as its directionality. Typically, the directionality of singular vectors of a matrix-power spectral density at a given frequency relates to the relative strength of the corresponding signal-components at the location of the measurement channels (sensors). Therefore, in order to quantify the performance of estimation algorithms, accurately detect changes in time series (events), localize the directionality of echo (e.g., in radar), etc., one needs physically meaningful metrics that weigh in relevant characteristics of power spectra. Thus, at the very least, metrics ought to be weakly continuous and allow the comparison of matrix-valued densities.

Several distance measures have been proposed and extensively used in applications and the literature, see e.g., [15], [16], [17], [18], [19], [1]. However, besides the fact that these fail to be metrics, most fail to be weakly continuous as well (e.g., the Itakura-Saito distance, etc.). In the present section, following the recipe outlined earlier, in (4-6), we

develop a weakly continuous metric between matrix-valued measures that can be thought of as a generalization of the Wasserstein metric. Our viewpoint herein differs from, yet it can be seen as complementing to the one in [1].

For a Hermitian matrix-valued function $\boldsymbol{f}$, let

$$\|\boldsymbol{f}\|_{\mathrm{Lip}} := \sup_{x \neq y} \frac{\|\boldsymbol{f}(x) - \boldsymbol{f}(y)\|}{d(x,y)}$$

where $d(x,y)$ is a metric on $\mathbb{I}$. Then $\|\boldsymbol{f}\|_{\mathrm{Lip}} \leq 1$ implies that

$$\boldsymbol{f}(x) - \boldsymbol{f}(y) \leq d(x,y) I$$

with I being the identity matrix in the sense of positive semi-definiteness. For any two matrix-valued power spectral measures $\boldsymbol{\mu}_1, \boldsymbol{\mu}_2$ we define:

$$\begin{aligned} &\mathrm{d}_{\mathrm{W}_1,\kappa}(\boldsymbol{\mu}_1, \boldsymbol{\mu}_2) = \\ &\sup_{\boldsymbol{f}} \{ \int_{\mathbb{I}} \mathrm{tr}(\boldsymbol{f}(d\boldsymbol{\mu}_1 - d\boldsymbol{\mu}_2) \mid \|\boldsymbol{f}\|_{\mathrm{Lip}} \leq 1, \|\boldsymbol{f}\| \leq \kappa \}. \end{aligned} \tag{7}$$

It is quite clear that this defines a metric and as we state next, it is also weakly continuous.

Proposition 2: Consider a sequence of matrix-valued power spectral measures $\{\boldsymbol{\mu}_k : k = 1, 2, \ldots\}$ and $\boldsymbol{\mu}$. Then

$$\mathrm{d}_{\mathrm{W}_1,\kappa}(\boldsymbol{\mu}_k, \boldsymbol{\mu}) \underset{k \to \infty}{\longrightarrow} 0 \tag{8}$$

if and only if for any continuous, bounded, Hermitian-valued function $\boldsymbol{f}$ on $\mathbb{I}$ the following holds

$$\mathrm{tr}\left(\int_{\mathbb{I}} \boldsymbol{f} d\boldsymbol{\mu}_k \right) \underset{k \to \infty}{\longrightarrow} \mathrm{tr}\left(\int_{\mathbb{I}} \boldsymbol{f} d\boldsymbol{\mu} \right). \tag{9}$$

Proof: See Appendix A. ∎

We provide an interpretation of (7) that draws a connection to optimal mass transport very much like in Section II. For this, we need the following expression for the total variation:

$$\|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_{\mathrm{TV}} := \int_{\mathbb{I}} \|d\boldsymbol{\mu}_1 - d\boldsymbol{\mu}_2\|_*$$

where $\|\cdot\|_*$ denotes the nuclear norm, i.e., the sum of singular values. We also need to the following matricial analog of the 1-Wasserstein metric between matrix-valued measures $\boldsymbol{\mu}_1$ and $\boldsymbol{\mu}_2$ with the same "total matricial mass" $\boldsymbol{\mu}_1(\mathbb{I}) = \boldsymbol{\mu}_2(\mathbb{I})$,

$$\begin{aligned} \mathrm{d}_{W_1}(\boldsymbol{\mu}_1, \boldsymbol{\mu}_2) = \inf_{\boldsymbol{m}} \Bigg\{ &\int_{\mathbb{I}\times\mathbb{I}} d(x,y) \|d\boldsymbol{m}(x,y)\|_* \mid \\ &\int_{y \in \mathbb{I}} d\boldsymbol{m}(x,y) = d\boldsymbol{\mu}_1(x), \int_{x \in \mathbb{I}} d\boldsymbol{m}(x,y) = d\boldsymbol{\mu}_2(y) \Bigg\}. \end{aligned}$$

We remark that, in the above, $d\boldsymbol{m}(x,y)$ needs not to be positive semidefinite. The optimization seeks a distribution for the nuclear norm of $d\boldsymbol{m}(x,y)$ which may now

be thought of actually as the "mass" transported from x to y. When $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$ are scalar-valued probability measure, clearly, $\mathrm{d}_{W_1}(\boldsymbol{\mu}, \boldsymbol{\nu})$ is the 1-Wasserstein metric. The following statement represents a generalization of (3).

Proposition 3: For two matrix-valued measures $\boldsymbol{\mu}_1$ and $\boldsymbol{\mu}_2$ on $\mathbb{I}$,

$$\mathrm{d}_{\mathrm{W}_{1,\kappa}}(\boldsymbol{\mu}_1, \boldsymbol{\mu}_2) := \inf_{\hat{\boldsymbol{\mu}}_1, \hat{\boldsymbol{\mu}}_2} \left\{ \mathrm{d}_{W_1}(\hat{\boldsymbol{\mu}}_1, \hat{\boldsymbol{\mu}}_2) + \kappa \sum_{k=1}^{2} \|\boldsymbol{\mu}_k - \hat{\boldsymbol{\mu}}_k\|_{\mathrm{TV}} \right\} \tag{10}$$

Proof: See Appendix B. ∎

V. Example

We highlight the characteristics of the proposed distance $\mathrm{d}_{\mathrm{W}_{1,\kappa}}$ with a numerical example. In this, we compare three matrix-valued densities $\boldsymbol{f}_0, \boldsymbol{f}_1, \boldsymbol{f}_2$ and compute the values assigned by the metric between them, for each pair. The nature and directionality of their spectral content is such that $\boldsymbol{f}_1$ can naturally be thought of as "sitting" in the middle of the other two. Thus, if a metric is to be intuitive, it ought to assign distances accordingly. We compare the relative values assigned by our metric as well as three alternatives. These are, d_{IS}, d_{TV} and the metric introduced in our earlier work [1]. It is seen that (7), as well as the metric in [1], are quite similar, while the other two, d_{IS} and d_{TV}, give relative distances that do not reflect the intuition suggested above.

The three chosen matricial densities are:

$$\boldsymbol{f}_0(\theta) = \begin{bmatrix} 1 & 0.4 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0.01 & 0 \\ 0 & \frac{1}{|a_0(e^{j\theta})|^2} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0.4 & 1 \end{bmatrix}$$

$$\boldsymbol{f}_1(\theta) = \begin{bmatrix} 1 & 0.5 \\ 0.5e^{j\theta} & 1 \end{bmatrix} \begin{bmatrix} \frac{1}{|a_1(e^{j\theta})|^2} & 0 \\ 0 & \frac{1}{|a_1(e^{j\theta})|^2} \end{bmatrix} \begin{bmatrix} 1 & 0.5e^{-j\theta} \\ 0.5 & 1 \end{bmatrix}$$

$$\boldsymbol{f}_2(\theta) = \begin{bmatrix} 1 & 0 \\ 0.4e^{j\theta} & 1 \end{bmatrix} \begin{bmatrix} \frac{1}{|a_2(e^{j\theta})|^2} & 0 \\ 0 & 0.01 \end{bmatrix} \begin{bmatrix} 1 & 0.4e^{-j\theta} \\ 0 & 1 \end{bmatrix}$$

where

$$\begin{aligned} a_0(z) &= (1 - 1.9\cos(\frac{\pi}{6})z + 0.95^2 z^2)(1 - 1.5\cos(\frac{\pi}{3})z + 0.75^2 z^2) \\ a_1(z) &= (1 - 1.9\cos(\frac{5\pi}{12})z + 0.95^2 z^2)(1 - 1.5\cos(\frac{\pi}{2})z + 0.75^2 z^2) \\ a_2(z) &= (1 - 1.9\cos(\frac{2\pi}{3})z + 0.95^2 z^2)(1 - 1.5\cos(\frac{5\pi}{8})z + 0.75^2 z^2) \end{aligned}$$

for $\theta \in [0, \pi]$. These are shown in Figure 1 and since, for each $i \in \{0, 1, 2\}$, $\boldsymbol{f}_i$ is a Hermitian-valued and positive semi-definite, our convention is to display the magnitude and phase of their entries as follows: $|\boldsymbol{f}_{i,(1,1)}|$, $|\boldsymbol{f}_{i,(1,2)}|$ $(= |\boldsymbol{f}_{i,(2,1)}|)$ and $|\boldsymbol{f}_{i,(2,2)}|$ are displayed in subplots (1,1), (1,2) and (2,2), respectively, and $\angle \boldsymbol{f}_{i,(1,2)}$ $(= -\angle \boldsymbol{f}_{i,(2,1)})$ is shown in subplot (2,1).

Peak power in $\boldsymbol{f}_0$ is at $\theta = \frac{\pi}{6}$ and most of it resides in the second channel, i.e., in $\boldsymbol{f}_{0,(2,2)}$. The power in $\boldsymbol{f}_1$ splits equality between the two channels and its peak is at $\theta = \frac{5\pi}{12}$. In $\boldsymbol{f}_2$, peak power is in the first channel and at around $\theta = \frac{2\pi}{3}$. Similarly, it is worth observing the characteristics and phase angles of the cross spectra. All in all, $\boldsymbol{f}_1$ appears to be "sitting in the middle" between $\boldsymbol{f}_0$ and $\boldsymbol{f}_2$. By comparing the relative distances we can now assess whether these reflect the above intuition, i.e., that $\boldsymbol{f}_1$ is in the middle between $\boldsymbol{f}_0$ and $\boldsymbol{f}_2$.

We compute and compare distances given by the generalized Itakura-Saito distance [15]

$$\mathrm{d}_{\mathrm{IS}}(\boldsymbol{f}_0, \boldsymbol{f}_1) := \int_0^\pi \mathrm{tr}\left(\boldsymbol{f}_0 \boldsymbol{f}_1^{-1} - \log(\boldsymbol{f}_0 \boldsymbol{f}_1^{-1}) - I\right) d\theta$$

the total variation $\mathrm{d}_{\mathrm{TV}}(\boldsymbol{f}_0, \boldsymbol{f}_1) := \|\boldsymbol{f}_0 - \boldsymbol{f}_1\|_{\mathrm{TV}}$ and the metric $\mathrm{d}_{\mathrm{T},\kappa}$ proposed in [1]. To this end, we sample the $\boldsymbol{f}_i$'s on a frequency grid with resolution $\Delta\theta = \frac{\pi}{35}$. To obtain $\mathrm{d}_{\mathrm{T},\kappa}$, we first scale so that traces are normalized to 1. The distances are given in Table I where we have chosen $\kappa = 1$ for both $\mathrm{d}_{\mathrm{W}_1,\kappa}$ and $\mathrm{d}_{\mathrm{T},\kappa}$.

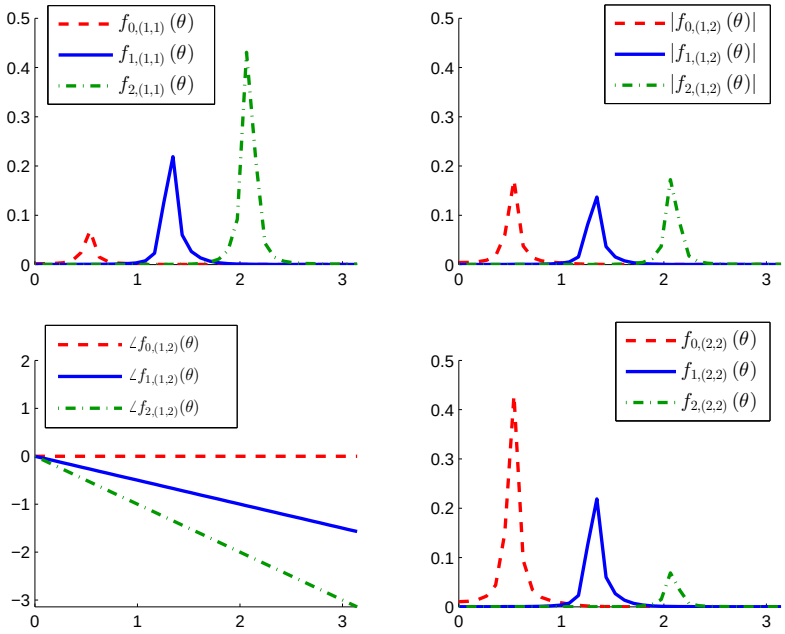

Fig. 1. Matrix-valued power spectra $\boldsymbol{f}_0$, $\boldsymbol{f}_1$ and $\boldsymbol{f}_2$ are shown in red dashed line, blue solid line and green dashdot line, respectively. Subplots (1,1), (1,2) and (2,2) show $\boldsymbol{f}_{i,(1,1)}, |\boldsymbol{f}_{i,(1,2)}|$ (same as $|\boldsymbol{f}_{i,(2,1)}|$) and $\boldsymbol{f}_{i,(2,2)}$. Subplot (2,1) shows $\angle(\boldsymbol{f}_{i,(1,2)})$ for $i \in \{0, 1, 2\}$.

We observe that d_{IS} gives a rather unintuitive result since we expect similar distances between the two pairs $\boldsymbol{f}_0, \boldsymbol{f}_1$ and $\boldsymbol{f}_1, \boldsymbol{f}_2$. The total variation does not differentiate the

	Distance between pairs of density functions		
	$\boldsymbol{f}_0, \boldsymbol{f}_1$	$\boldsymbol{f}_1, \boldsymbol{f}_2$	$\boldsymbol{f}_0, \boldsymbol{f}_2$
$\mathrm{d_{IS}}$	3.44×10^3	5.36×10^4	9.27×10^4
$\mathrm{d_{TV}}$	1.95	1.96	2.00
$\mathrm{d_{W_1,\kappa}}$	1.37	1.65	2.29
$\mathrm{d_{T,\kappa}}$	1.01	1.09	2.05

TABLE I
DISTANCES BETWEEN DENSITY FUNCTIONS.

three pairs since all three are found equally close to each other. Using $\mathrm{d_{T,1}}$, $\mathrm{d_{T,1}}(\boldsymbol{f}_0, \boldsymbol{f}_1)$ is very close to $\mathrm{d_{T,1}}(\boldsymbol{f}_1, \boldsymbol{f}_2)$ and they are both at nearly half compared to $\mathrm{d_{T,1}}(\boldsymbol{f}_0, \boldsymbol{f}_2)$. Then $\mathrm{d_{W_1,1}}$ is quite similar to $\mathrm{d_{T,1}}$. These comparisons suggest that $\mathrm{d_{W_1,\kappa}}$ as well as $\mathrm{d_{T,1}}$, both reflect quite closely the intuition and what one would expect after inspecting the relative distribution of power and directionality in all three densities.

VI. CONCLUDING REMARKS

The main thesis of this paper is that it is natural to quantify distances by evaluating measures against suitable families of test functions. After explaining the general recipe on representative metrics on scalar densities, including the one proposed in our earlier work [10], we expand on possible directions that lead to metrics for matrix-valued measures and density functions. We discuss in detail one so-derived weakly continuous metric which can be thought of as a natural generalization of the 1-Wasserstein metric between matrix-valued densities. Comparison with alternatives as well as a similar metric in [1] is explained on an academic example. A key property of the metric introduced herein is the weak continuity which is not shared by any of the alternatives.

APPENDIX A
PROOF OF PROPOSITION 2

First, it is clearly (9) implies (8). We show that (8) also implies (9). The proof below follows that in [7, page 216] for the scalar-valued Wasserstein metric.

We notice that if (8) holds, then (9) hold for any Lipschitz and bounded $\boldsymbol{f}$. To see this, for any Lipschitz and bounded $\boldsymbol{f}$, it can be scaled so that $\boldsymbol{f}/\max\{\|\boldsymbol{f}\|_{\mathrm{Lip}}, \frac{\max\|\boldsymbol{f}\|}{\kappa}\}$ is a feasible element in (7). Then, the proof requires showing that any element in the set of continuous, bounded Hermitian matrix-valued function $C_b(\mathbb{H}^{\ell\times\ell}, \mathbb{I})$ can be approached by sequences of Lipschitz and bounded ones from above and below respectively.

We need the following expression for suitable $\inf$ and the $\sup$ of an $\boldsymbol{f} \in C_b(\mathbb{H}^{\ell\times\ell}, \mathbb{I})$:

$$\begin{aligned}\inf_{x\in\mathbb{I}} \boldsymbol{f}(x) &:= \arg\sup_{\boldsymbol{m}} \{\mathrm{tr}(\boldsymbol{m}) \mid \boldsymbol{f}(x) - \boldsymbol{m} \geq 0\}\\ \sup_{x\in\mathbb{I}} \boldsymbol{f}(x) &:= \arg\inf_{\boldsymbol{m}} \{\mathrm{tr}(\boldsymbol{m}) \mid \boldsymbol{m} - \boldsymbol{f}(x) \geq 0\}\end{aligned}$$

with ordering in the sense of positive semi-definiteness. For $n \in \mathbb{N}$, we denote

$$\begin{aligned} \boldsymbol{f}_{\mathrm{low},n}(x) &:= \inf_{y\in\mathbb{I}} \{\boldsymbol{f}(y) + nd(x,y)I\} \\ \boldsymbol{f}_{\mathrm{up},n}(x) &:= \sup_{y\in\mathbb{I}} \{\boldsymbol{f}(y) - nd(x,y)I\}. \end{aligned}$$

Then, $\boldsymbol{f}_{\mathrm{low},n}(x) \le \boldsymbol{f}(x) \le \boldsymbol{f}_{\mathrm{up},n}(x)$. For any fixed x, $\{\boldsymbol{f}_{\mathrm{low},n}(x) : n \in \mathbb{N}\}$ and $\{\boldsymbol{f}_{\mathrm{up},n}(x) : n \in \mathbb{N}\}$ are increasing sequences and decreasing sequences, respectively, and they both converge to $\boldsymbol{f}(x)$. It is also important to note that $\|\boldsymbol{f}_{\mathrm{low},n}\|_{\mathrm{Lip}} \le n$ and $\|\boldsymbol{f}_{\mathrm{up},n}\|_{\mathrm{Lip}} \le n$. To see this, for any $x, y \in [0,1]$, we have

$$\begin{aligned} &\boldsymbol{f}_{\mathrm{low},n}(x) - \boldsymbol{f}_{\mathrm{low},n}(y) \\ &= \inf_z \left\{\boldsymbol{f}(z) + nd(x,z)I\right\} - \inf_{\hat{z}} \left\{\boldsymbol{f}(\hat{z}) + nd(y,\hat{z})I\right\} \\ &\le \sup_z \left\{(\boldsymbol{f}(z) + nd(x,z)I) - (\boldsymbol{f}(z) + nd(y,z)I)\right\} \\ &= \sup_z n(d(x,z) - d(y,z))I \\ &\le nd(x,y)I. \end{aligned}$$

Thus $\|\boldsymbol{f}_{\mathrm{low},n}\|_{\mathrm{Lip}} \le n$, and $\|\boldsymbol{f}_{\mathrm{up},n}\|_{\mathrm{Lip}} \le n$ can be proved in a similar manner. Then we have

$$\begin{aligned} &\limsup_{k\to\infty} \mathrm{tr}\left(\int_0^1 \boldsymbol{f} d\boldsymbol{\mu}_k\right) \\ &\le \liminf_{n\to\infty} \limsup_{k\to\infty} \mathrm{tr}\left(\int_0^1 \boldsymbol{f}_{\mathrm{up},n} d\boldsymbol{\mu}_k\right) \\ &= \liminf_{n\to\infty} \mathrm{tr}\left(\int_0^1 \boldsymbol{f}_{\mathrm{up},n} d\boldsymbol{\mu}\right) \\ &= \mathrm{tr}\left(\int_0^1 \boldsymbol{f} d\boldsymbol{\mu}\right). \end{aligned}$$

Similarly, using the sequence $\{\boldsymbol{f}_{\mathrm{low},n}(x) : n \in \mathbb{N}\}$ we can also show that

$$\liminf_{k\to\infty} \mathrm{tr}\left(\int_0^1 \boldsymbol{f} d\boldsymbol{\mu}_k\right) \ge \mathrm{tr}\left(\int_0^1 \boldsymbol{f} d\boldsymbol{\mu}\right).$$

This completes the proof.

APPENDIX B
PROOF OF PROPOSITION 3

We show that (7) is dual of (10): rewrite (10) as

$$\begin{aligned} \inf_{\boldsymbol{m},\hat{\boldsymbol{\mu}},\hat{\boldsymbol{\nu}}} \Bigg\{ &\int_{\mathbb{I}\times\mathbb{I}} d(x,y)\|\boldsymbol{m}(x,y)\|_* + \kappa \sum_{k=1}^{2} \|\boldsymbol{\mu}_k - \hat{\boldsymbol{\mu}}_k\|_{\mathrm{TV}} \quad | \\ &\int_{y\in\mathbb{I}} d\boldsymbol{m}(x,y) = d\hat{\boldsymbol{\mu}}_1(x), \quad \int_{x\in\mathbb{I}} d\boldsymbol{m}(x,y) = d\hat{\boldsymbol{\mu}}_2(y) \Bigg\}. \end{aligned} \tag{11}$$

To derive the dual formulation, we need to rewrite the nuclear norm as follows

$$\|d\boldsymbol{m}\|_* = \max_{\|\boldsymbol{w}\|\leq 1} \operatorname{tr}(\boldsymbol{w} d\boldsymbol{m})$$
$$\|d\boldsymbol{\mu}_k - \hat{d}\boldsymbol{\mu}_k\|_* = \max_{\|\boldsymbol{\lambda}_k\|\leq 1} \operatorname{tr}(\boldsymbol{\lambda}_k(d\boldsymbol{\mu}_k - d\boldsymbol{\mu}_k)).$$

Let $\boldsymbol{\phi}_1, \boldsymbol{\phi}_2$ be the Lagrange multipliers for the two constraints. Using the Lagrange multiplier method, (11) equals the following

$$\begin{aligned}
\inf_{\boldsymbol{m},\hat{\boldsymbol{\mu}}_1,\hat{\boldsymbol{\mu}}_2} \sup_{\substack{\boldsymbol{\phi}_1,\boldsymbol{\phi}_2 \\ \|\boldsymbol{w}\|,\|\boldsymbol{\lambda}_0\|,\|\boldsymbol{\lambda}_1\|\leq 1}} \Bigg\{ & \int \operatorname{tr}(\boldsymbol{\phi}_1(x) - \kappa\boldsymbol{\lambda}_0(x))\hat{d}\boldsymbol{\mu}_1(x) \\
& + \int \operatorname{tr}(\boldsymbol{\phi}_2(y) - \kappa\boldsymbol{\lambda}_2(y))\hat{d}\boldsymbol{\mu}_2(y) \\
& + \int \operatorname{tr}\left(d(x,y)\boldsymbol{w}(x,y) - \boldsymbol{\phi}_1(x) - \boldsymbol{\phi}_2(y)\right) d\boldsymbol{m}(x,y) \\
& + \kappa \int \operatorname{tr}\Big(\boldsymbol{\lambda}_1(x)d\boldsymbol{\mu}_1(x) + \boldsymbol{\lambda}_2(x)\boldsymbol{\mu}_2(x)\Big)\Bigg\}. \qquad (12a)
\end{aligned}$$

Since this optimization problem is convex, the optimal value does not change by switching the $\inf$ and $\sup$ in (12a). The optimal assignement for $\boldsymbol{w}, \boldsymbol{\phi}_k, \boldsymbol{\lambda}_k$ must satisfy

$$\begin{aligned}
d(x,y)\boldsymbol{w}(x,y) - \boldsymbol{\phi}_1(x) - \boldsymbol{\phi}_2(y) &= 0, \qquad (13a) \\
\boldsymbol{\phi}_1(x) - \kappa\boldsymbol{\lambda}_0(x) &= 0, \qquad (13b) \\
\boldsymbol{\phi}_2(y) - \kappa\boldsymbol{\lambda}_1(y) &= 0. \qquad (13c)
\end{aligned}$$

Setting $x = y$ in (13a), we obtain $\boldsymbol{\phi}_1 = -\boldsymbol{\phi}_2 =: \boldsymbol{\phi}$. Thus

$$\boldsymbol{\phi}(x) - \boldsymbol{\phi}(y) \leq d(x,y)I, \text{ and } \|\boldsymbol{\phi}\| \leq \kappa.$$

By substituting these conditions to (12a), we derive (7).

REFERENCES

[1] L. Ning, T. Georgiou, and A. Tannenbaum, "Matrix-valued Monge-Kantorovich optimal mass transport," *IEEE Transactions on Automatic Control*, to appear, 2015.

[2] L. Ning, "Matrix-valued optimal mass transportation and its applications," Ph.D. dissertation, University of Minnesota, December 2013.

[3] V. Zolotarev, "Probability metrics," *Theory of Probability & Its Applications*, vol. 28, no. 2, pp. 278–302, 1984.

[4] A. Müller, "Integral probability metrics and their generating classes of functions," *Advances in Applied Probability*, pp. 429–443, 1997.

[5] S. T. Rachev, L. Klebanov, S. V. Stoyanov, and F. Fabozzi, *The methods of distances in the theory of probability and statistics*. Springer, 2013.

[6] S. T. Rachev, *Probability metrics and the stability of stochastic models*. Wiley Chichester, 1991, vol. 334.

[7] C. Villani, *Topics in optimal transportation*. American Mathematical Society, 2003, vol. 58.

[8] X. Jiang, Z. Luo, and T. Georgiou, "Geometric methods for spectral analysis," *IEEE Transactions on Signal Processing*, vol. 60, no. 3, pp. 1064–1074, 2012.

[9] L. Ning, X. Jiang, and T. Georgiou, "On the geometry of covariance matrices," *IEEE Signal Processing Letters*, vol. 20, no. 8, pp. 787–790, 2013.

[10] T. Georgiou, J. Karlsson, and M. Takyar, "Metrics for power spectra: an axiomatic approach," *IEEE Transactions on Signal Processing*, vol. 57, no. 3, pp. 859–867, 2009.

[11] J. Karlsson and T. Georgiou, "Uncertainty bounds for spectral estimation," *IEEE Transactions on Automatic Control*, vol. 58, no. 7, pp. 1659–1673, 2013.

[12] A. Connes, *Noncommutative geometry*. Academic press, 1995.

[13] M. Rieffel, "Metrics on state spaces," *Doc. Math., J. DMV*, vol. 4, pp. 559–600, 1999.

[14] F. D'Andrea and P. Martinetti, "A view on optimal transport from noncommutative geometry," *SIGMA*, vol. 6, no. 057, p. 24, 2010.

[15] A. Ferrante, C. Masiero, and M. Pavon, "Time and spectral domain relative entropy: A new approach to multivariate spectral estimation," *IEEE Transactions on Automatic Control*, vol. 57, no. 10, pp. 2561–2575, 2012.

[16] A. Ferrante, M. Pavon, and F. Ramponi, "Hellinger versus Kullback–Leibler multivariable spectrum approximation," *IEEE Transactions on Automatic Control*, vol. 53, no. 4, pp. 954–967, 2008.

[17] A. Ferrante, M. Pavon, and M. Zorzi, "A maximum entropy enhancement for a family of high-resolution spectral estimators," *IEEE Transactions on Automatic Control*, vol. 57, no. 2, pp. 318–329, 2012.

[18] B. Afsari and R. Vidal, "The alignment distance on space of linear dynamical systems," in *IEEE Conference on Decision and Control*, 2013.

[19] X. Jiang, L. Ning, and T. Georgiou, "Distances and Riemannian metrics for multivariate spectral densities," *IEEE Transactions on Automatic Control*, vol. 57, no. 7, pp. 1723–1735, 2012.